Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Learning & Adaptive Systems

The Human Kernel Continued

Mojmír Mutný Andreas Krause
Computer Science, ETH Zürich, 8092 Switzerland

1 Motivation

In Bayesian optimization [Snoek et al., 2012], experimental design [Chaloner and Verdinelli, 1995], and general machine learning [Rasmussen and Williams, 2006], kernel methods and related probabilistic techniques called Gaussian processes have proven to be versatile and successful tools in practice. A common challenge is the selection of an appropriate kernel, which influences both the difficulty of the problem and its generalization properties.

To define such kernels, various methods are employed to learn from data and/or previous examples of similar problems. The gold standard occurs when a user possesses sufficient understanding of the problem and can define the relevant kernel themselves. However, in many cases, this task is quite challenging and seldom accomplished, leading to the reliance on uninformative kernels.

Despite this, there are experts who possess a deep understanding of the problem but struggle to encapsulate their knowledge into a kernel description. In other words, they know the kernels but are unaware of how to express them.

2 Challenge

In this project, we aim to design an interactive mechanism that sequentially interacts with an expert to determine the appropriate kernel by posing a series of questions. This approach builds upon the seminal work on human kernels [Wilson et al., 2015] and extends it based on the knowledge and experience accumulated within our group.

3 Background

We are seeking a highly motivated candidate with a strong background in machine learning and/or statistics. For more information, please contact Mojmir Mutny at mojmir.mutny@inf.ethz.ch. This project will be conducted in collaboration with a US university as part of an international partnership.

References

[Chaloner and Verdinelli, 1995] Chaloner, K. and Verdinelli, I. (1995). Bayesian experimental design: A review. *Statist. Sci.*, 10(3):273–304.

[Rasmussen and Williams, 2006] Rasmussen, C. and Williams, C. (2006). Gaussian processes for machine learning, vol. 1. *The MIT Press, Cambridge, doi*, 10:S0129065704001899.

[Snoek et al., 2012] Snoek, J., Larochelle, H., and Adams, R. P. (2012). Practical Bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pages 2951–2959.

[Wilson et al., 2015] Wilson, A. G., Dann, C., and Nickisch, H. (2015). Thoughts on massively scalable Gaussian processes. *arXiv preprint arXiv:1511.01870*.